

PProx

Simon Da Silva

Video streaming platforms (such as YouTube, Vimeo or Dailymotion) often feature recommendations for similar content to end-users. They rely on these recommendations to retain users on their website or application. To do so, they establish a user profile based on the watching history. This leads to threats to privacy as (i) service providers gather private data on each user, (ii) an adversary can intercept recommendations and deduce private information about the user, or (iii) malicious platform providers can display targeted ads tailored to a specific user instead of a generic profile to generate income.

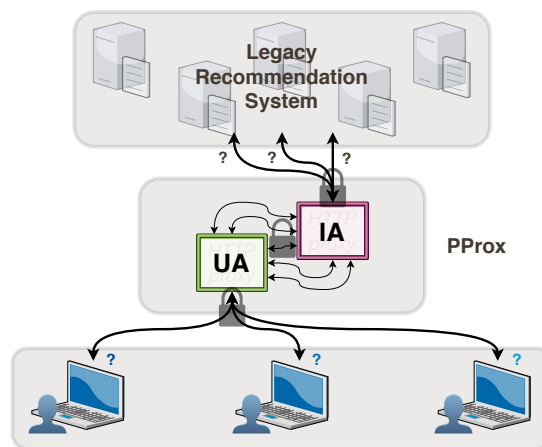
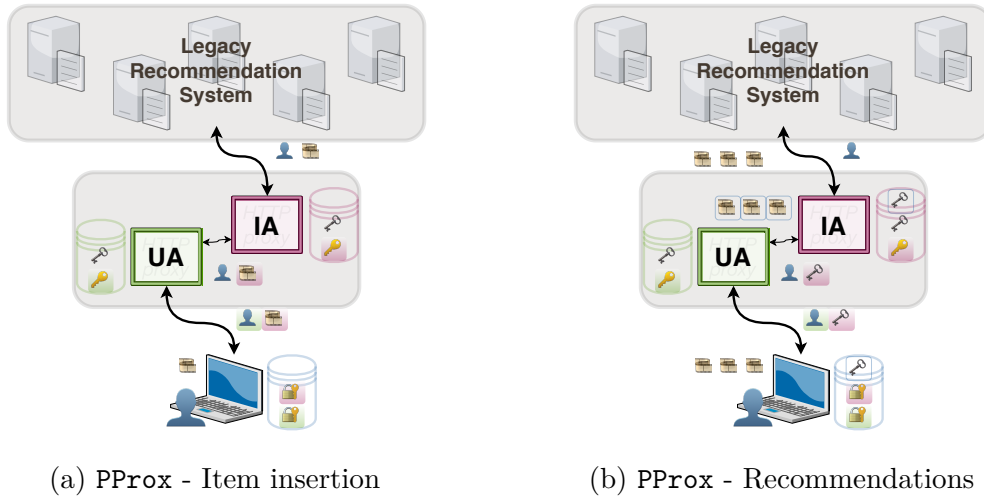


Figure 1: PProx overview

PProx is an efficient and easily-deployable solution for privacy preservation in video streaming recommendation engines. PProx does not impact recommendations accuracy, supports arbitrary recommendation algorithms, and has minimal deployment requirements. To tackle these issues, PProx introduces a privacy-preserving proxy service, standing between users and any off-the-shelf recommender system. This proxy service intercepts feedback insertions and requests for recommendations.

It pseudonymizes on the fly the user and items identifiers, hides links between the two, and shuffles all requests (see Figure 1). This guarantees unlinkability between clients and the items they access or receive as recommendations. The security guarantees of PProx hold even in the presence of a powerful attacker able to use recently-documented *side-channel* attacks on SGX enclaves to corrupt one of the proxies, and observing all network traffic in the cloud. PProx proxies elastically scale over a fleet of Intel SGX-enabled machines.



We integrated PProx with the Harness universal recommendation engine and evaluated it on a 27-node cluster. Our results indicate its ability to withstand a high number of requests with low end-to-end latency, scaling up to match the workload of recommendations. The typical latency overhead is below 100ms (compared to several seconds for similar systems), and PProx only requires 30% to 50% additional nodes to provide the recommendation service.